64. The method of claim 61, further comprising:

for each policy in the training set, providing a random value for the previously audited variable, and applying the derived variables and the random value of the previously audited variable to the predictive model; and

for each policy in the training set, providing an actual value for the previously audited variable indicating whether the policy was previously audited for the scoring period, and applying the derived variables and the actual value of the previously audited variable to calibrate the scores produced by the predictive model.

<u>REMARKS</u>

Claims 1-64 were presented for examination, and were rejected. Claims 42-44 are cancelled. Claims 2, 5-8, 14 and 58-61 are amended to correct various typographical errors. No claim is being amended for any reason related to patentability, except insofar as it may have been rejected by the Examiner because of a typographical error. No amendment is intended to narrow the scope of any claim.

The Examiner required a restriction under 35 U.S.C. 121 between Group I (claims 1-41 and 45-64) and Group II (claims 42-44). As agreed in a telephone conversation between the Examiner and Attorney Robert Sachs, Group I was elected without traverse. Claims 42-44 have therefore been withdrawn from consideration, and are hereby cancelled without prejudice.

The specification has been amended to correct occasional typographical errors. In addition, the Abstract has been amended to overcome the Examiner's length objection, and is now less than 150 words.

The Examiner issued a requirement for information under 37 C.F.R. 1.105. In response, Applicants are enclosing an HNC brochure describing HNC's Vericomp Claimant product. To the best of the Applicants' knowledge, Vericomp Claimant is the product referred to in the published document cited by the Examiner in the Requirement for Information. An Information Disclosure Statement citing the enclosed reference is also included.

The Examiner objected to claim 14 because of a typographical error. Claim 14 is amended to remove the typographical error.

The Examiner rejected claim 6 under 35 U.S.C. 112, second paragraph. Claim 6 now depends from claim 5, and thus provides antecedent basis for "the scoring period."

The Examiner rejected claims 1-11, 17-41 and 49-52 under 35 USC 103(a) as being unpatentable over Gopinathan et al. (Gopinathan) in view of Fischthal and Downs. Applicants respectfully traverse these rejections.

Claim 1 recites:

> A method for detecting misrepresentation of policy related information provided to an insurer by a policyholder where the information is used by the insurer in determining an amount of premium to be paid for insurance coverage provided to the policyholder, the method comprising:
>
> selecting a plurality of insurance policies to process with a predictive model;
>
> for each selected policy, deriving variables from policy related information provided by the policyholder in connection with the selected policy; and
>
> for each selected policy, applying the derived variables of the policy to the predictive model to generate a model score

> indicating the relative likelihood of misrepresented information provided by the policyholder or an expected adjustment of the premium on the policy.

In determining appropriate premiums to charge their customers, insurers rely in large part on information provided by the customers themselves. Thus, if a customer provides false or misleading information, the misrepresentation can result in a premium that is too low. The claimed invention detects such misrepresentations by using a predictive model to analyze variables derived from selected policies, and generating a score indicating the relative likelihood of misrepresentation in the policy information provided.

Gopinathan does not disclose, teach or suggest the claimed invention, either alone or in combination with the Fischthal or Downs references. Gopinathan relates generally to using a predictive model to detect fraudulent use of customer accounts and account numbers, such as in credit card transactions (col. 1, lines 12-15). Although both Gopinathan and the claimed invention disclose using a predictive model to determine a score, the two are patentably distinct. For example, in Gopinathan, credit card users are profiled, and each transaction is scored for fraud in view of that profile. In contrast, the claimed invention <u>does not use profiles</u>, and indeed <u>does not even score transactions</u>. Instead, the claimed invention scores information submitted periodically, such as payroll reports in the case of workers' compensation. To that end, while Gopinathan teaches instantly determining a score for a discrete transaction, based on a profile and without reference to a database to derive profiles, the claimed invention derives variables from policy information, typically provided some time ago, as well as claim information, and does so without any profiles by referring to a database.

Thus, Gopinathan uses profiles to score instantaneous, discrete transactions in order to assess whether the transaction being scored is fraudulent. The claimed

invention uses no profiles, and derives variables from data provided by a policyholder in order to determine whether the policyholder should be audited. There is no transaction in the claimed invention, and indeed no analogy between assessing whether a credit card transaction is fraudulent, and whether an insurance policy has been obtained fraudulently. In the former case, fraud typically arises because a valid, lawfully obtained credit card is being used by a person not authorized to use it. In the latter case, using workers' compensation as an example, employer fraud arises when the employer lies about the nature of the business, the number of employees, the payroll amounts, etc., in order to reduce the premium. The Gopinathan model of using profiles to score transactions does not lend itself to solving this problem, and indeed Gopinathan does not teach, disclose or suggest the claimed invention.

The addition of Fischtal does not cure the defects of Gopinathan. Fischthal uses data-driven clustering to define groups (classes) and then trains a neural network <u>on each of the classes from the clustering.</u> For example, if his clustering technique creates 10 classes, he builds 10 neural networks. That technique is a form of "segmentation." <u>The claimed invention does not do any segmentation</u>. The claimed invention uses clusters (industry-defined or data-driven) to help derive variables (e.g. distribution of payroll in this payroll report compared with its peers). A single neural network is then trained on all the data. While Fischthal handles different groups by segmenting and building multiple neural networks, the claimed invention handles different groups more elegantly by creating variables that enable a policy to be handled correctly (no matter what group it is in) all within a single neural network. As Fischthal points out, his method works with "very large amounts of data." One drawback to segmentation (by any method) is that with ten classes, one needs enough data to build ten neural networks. Without segmenting, only enough data to build one neural network is needed, but in order to do that, a way is needed to handle (within one model) the fact

that the data encompasses several groups (e.g. policies from several industries -- a construction firm is not the same thing as a law firm and their payroll reports would be expected to look different). Fischthal's technique is understood to have been developed while working on IRS data, and because of the very large amount of data, his technique lends itself to multiple neural networks. However, such a system does not lend itself to most fraud detection applications, including that of the claimed invention. Thus, Fischtal actually teaches away from the claimed invention, and does not, either alone or in combination with Gopinathan, teach, suggest or disclose the claimed invention.

The addition of the Downs article does not cure the defects of Gopinathan and Fischtal. The Downs article merely discloses the existence of the health care fraud problem. Morever, the article is primarily concerned with fraud in health care claims, not in the obtaining of lower premiums through misrepresentation. Thus, the Downs article does not teach, suggest or disclose claim 1.

Thus, neither Gopinathan, Fischtal, nor the Downs article, alone or in combination, teach, suggest or disclose claim 1. Independent claim 49 is patentable over the cited references for reasons analogous to those above with respect to claim 1. Dependent claims 2-41, 45-48 and 50-52 depend from patentable claims 1 and 49, and derive patentability from the independent claims from which they depend, in addition to reciting their own patentable features. Thus, claims 1-41 and 45-52 are patentable over the cited references and should be allowed.

The Examiner rejected claims 53-64 under 35 U.S.C. 103(a) as being unpatentable over Gopinathan in view of Fischthal, Downs, and Hann. However, claims 53-64 are patentable over Gopinathan in combination with Fischthal and Downs for the reasons discussed above with respect to claims 1 and 49. Furthermore, the addition of Hann does not cure the defects of the remaining references. Hann discloses the existence of

systems designed to detect claims fraud, such as HNC's VeriComp Claimant Fraud and Abuse Detection System (VCC). However, as is the case with Gopinathan, VCC and other systems are designed to use profiles to score transactions (including claims). The claimed invention does not use profiles, and does not score transactions, as discussed above. Therefore, claims 53-64 are patenable over the cited references and the rejection should be withdrawn.

In light of these Remarks, the Examiner is asked to issue a Notice of Allowance allowing all claims now pending, claims 1-41 and 45-64. If any issues remain outstanding prior to allowance, the Examiner is requested to contact the undersigned attorney so that they may be expeditiously resolved.

Respectfully submitted,
Ho Ming Luk, Pamela E. Coates,
Arati S. Deo, Sean M. Downs,
Benjamin A. Friesen,
Craig A. Nies, and Anu K. Pathria

Date: 29 July 2002.

By: DRB

Daniel R. Brownstone, Reg. No. 46,581
FENWICK & WEST LLP
Two Palo Alto Square
Palo Alto, CA 94306
Tel: (415) 875-2358
Fax: (415) 281-1350

Version with Markings to Show Changes Made

IN THE SUMMARY

Paragraph beginning at line 11 of page 9:

Referring to Fig. 2 there is now shown a conceptual diagram of how the predictive model detects premium fraud. The collection of insurance policies upon which the predictive model is developed [from] form a complex multi-dimensional "policy space" 201, which contains all of the policies that will be evaluated by the predictive model. Each policy is described by many policy variables. These policy variables generally fall into three categories of variables: over-time policy variables 203, peer group variables 205, and internal policy variables 207. It is this collection of policy variables that describes each policy in the policy space 201. In general, many of these variables may be understood as measures of the amount, distribution, or nature of the activities or characteristics of the policyholder and its claimants as indicators of premium fraud risk.

IN THE DETAILED DESCRIPTION

Paragraph beginning at line 20 of page 42:

4. Policies with an officer who is currently or was recently an officer on a different policy and where the new policy has a lower experience modification rate than the previous policy. The logic here attempts to identify policies that may be evading high experience modification rates by closing the company and re-opening it under a new name.

5. Policies that have a class code on a claim for which no premium was reported at the time the claim was opened. The logic here is similar to the first rule,

except in this case the job class code is listed on the payroll report but no payroll is reported in that class code. This may imply that the employer is misrepresenting the job classifications of their payroll in order to lower their premium.

Each rule in the rule-based analysis 620 flags any policies that violate the rule. These flags can be used to create lists of violators, which are useful complements to the scores from the predictive model 622. As noted above, in a workers' compensation implementation, policies with zero payroll are not scored by the predictive model 622, so without the rule-based analysis, suspicious policies in that group would not be evaluated. While the exclusion of such policies from the predictive model 622 is appropriate, it may still be possible to identify suspicious policies in this group, as the above rules demonstrate. Thus, the rule-based analysis 620 provides such analysis, bringing any problem policies with zero payroll to the attention of auditors. The rule-based analysis can also provide valuable additional analysis for policies that are scored by the predictive model 622. For example, a policy with a class code on a claim that is not on the policy might be scored by the predictive model 622, but if nothing else about that policy looks suspicious, it may not score high. The rule-based analysis 620 however would flag such a policy as having a clear-cut, specific problem that is independent of how suspicious the policy looks more generally.

Paragraph beginning at line 2 of page 68:

A randomly selected portion (*e.g.*, 20-30%) of the model development dataset is held out 909 from model training. This hold-out set is referred to as the "test" data 908b and is used to test the model that is trained on the remaining dataset 908a portion of the dataset 906a. Evaluation of the hold-out data ensures that the predictive model 622 does not over-fit the training data 908a. Also, the test data can be used to estimate [of] the production performance of the model (indeed, of the entire system).

IN THE ABSTRACT

Detection of insurance premium fraud [due to the misrepresentation of policy related information by the policyholder] is provided by a predictive model, which uses derived variables to assess the likelihood of [premium] fraud for each policy. [The variables are derived from data about the policy and similar peer policies. The variables capture selected information about the policy, changes over time in the policy's behavior or characteristics, and compare the policy with its peers.] The predictive model produces a score [for the policy], which is a measure of the likelihood of [a misrepresentation of policy information, and thus] premium fraud or abuse. [It also provides information on the factors that most strongly contribute to the score.] The predictive model is included in a system that accepts policies to be considered for scoring, selects which policies are appropriate for scoring, stores data about the policies in a database, uses the data to derive variables for the model, and processes and outputs the model scores and related information. A rule-based analysis, which detects specific inconsistencies in the data that are indicative of premium fraud, may also be part of the system. [The rule-based analysis may analyze policies even if they were rejected for scoring by the predictive model. Policies may be presented to the system automatically or interactively.] The model scores and red-flag indicators from the rule-based analysis may be further processed to provide customized output for users. [Insurers use the results to identify suspects of premium fraud and to prioritize audits and other investigations.]

IN THE CLAIMS

2. (Amended) The method of claim 1, further comprising:

collecting training data including a plurality of insurance policies having misrepresented information and a plurality of policies not having misrepresented information;

developing the predictive model from the training data; and

storing the predictive model.

5. (Amended) The method of claim 1, wherein selecting a plurality of insurance policies further comprises:

for each policy, automatically determining start and end dates of a scoring period [in] over which [the determination of whether] misrepresented policy information is to be [determined] detected.

6. (Amended) The method of claim [1] 5, [further comprising determining] wherein the start and end dates of the scoring period include a period for which the policy has consistent and complete data.

7. (Amended) The method of claim [6] 5, further comprising:

responsive to a policy not having consistent or complete data in the scoring period, defining an exclusion code providing a reason that the policy was not selected.

8. (Amended) The method of claim [6] 5, wherein the insurance policies are workers' compensation insurance policies, and automatically determining start and end dates of the scoring period further comprises:

defining the start and end dates such that all audit adjustments are contained between the start and end dates.

14. (Amended) The method of claim 12, wherein deriving variables for the policy which compare the policy to other policies in its peer group(s) further comprises deriving variables that compare either at least one characteristic [o] of the policy with at least one corresponding characteristic of the policies in its peer group(s).

58. (Amended) The method of claim 53, further comprising:

establishing a plurality of reason codes, each reason code providing an explanation for a policy receiving a score; and

establishing for each of a number of reason codes, at least one audit action to be taken in response to a policy having a score which produces the reason code.

59. (Amended) A method for processing insurance policies suspected of premium fraud, the method comprising:

scoring each of a plurality of insurance policies with a predictive model to generate for each policy a score indicating a relative likelihood of premium fraud;

ranking the scored policies according to the scores;

selecting for a desk audit those policies having a score exceeding a first threshold score; and

selecting for a field audit those policies having a score exceeding a second threshold score, wherein the second threshold score is greater than the first threshold score.

60. (Amended) A method for processing insurance policies suspected of premium fraud, the method comprising:

scoring each of a plurality of insurance policies with a predictive model to generate for each policy a score indicating a relative likelihood of premium fraud;

determining for each scored policy an expected premium adjustment;

ranking the scored policies according to their expected premium adjustments;

selecting for a desk audit those policies having an expected premium adjustment exceeding a first threshold amount; and

selecting for a field audit those policies having a expected premium adjustment exceeding a second threshold amount, wherein the second threshold amount is greater than the first threshold amount.

61. (Amended) A method of developing a predictive model of insurance premium fraud, the method comprising:

collecting from at least one insurance company policy information for a plurality of insurance policies;

determining for each policy a scoring period for scoring the policy;

selecting a training set of policies;

deriving for each policy in the training set a plurality of variables from the policy information and from other information relevant to policy premiums;

applying the derived variables to an untrained predictive model to train the predictive model to produce a measure with respect to whether the policies are fraudulent or non-fraudulent during their respective scoring periods ; and

selecting a subset of the derived variables for [the] using in the predictive model, which variables significantly contribute to a prediction of whether a policy is fraudulent during its scoring period.